arXiv:1410.0412v2 [cs.DC] 23 Dec 2015

Modeling and analyzing performance for highly optimized propagation steps of the lattice Boltzmann method on sparse lattices

M. Wittmann*, T. Zeiser*, G. Hager*, and G. Wellein**

*Erlangen Regional Computing Center, University of Erlangen-Nuremberg, Germany

**Department of Computer Science, University of Erlangen-Nuremberg, Germany

December 24, 2015

Computational fluid dynamics (CFD) requires a vast amount of compute cycles on contemporary large-scale parallel computers. Hence, performance optimization is a pivotal activity in this field of computational science. Not only does it reduce the time to solution, but it also allows to minimize the energy consumption. In this work we study performance optimizations for an MPI-parallel lattice Boltzmann-based flow solver that uses a sparse lattice representation with indirect addressing. First we describe how this indirect addressing can be minimized in order to increase the single-core and chip-level performance. Second, the communication overhead is reduced via appropriate partitioning, but maintaining the single core performance improvements. Both optimizations allow to run the solver at an operating point with minimal energy consumption.

1 Introduction and related work

Performance optimization for computational fluid dynamics (CFD) should start at the core and chip level. In most CFD algorithms the relevant bottleneck is the main memory bandwidth. A high-quality implementation therefore should exhaust the available bandwidth as well as cause the least amount of memory data traffic per work unit.

In case of lattice Boltzmann methods (LBM) the decisive quantity is called *loop balance* (B_l) and is measured in bytes per fluid lattice site update ($\frac{\text{B}}{\text{FLUP}}$) [16]. There are several propagation step variants for LBM that achieve lowest data traffic per FLUP: the two-grid one-step algorithm with non-temporal stores [12], Bailey et. al's AA-pattern [1], and Geier's Eso-Twist [7, 9]. The latter two both work with a single grid only and arrange the processing order in a clever way to work around data dependencies. The first two variants were successfully implemented in the fluid flow solver framework ILBDC [15], which relies on a sparse lattice structure of the simulation domain [19]. In contrast to a full array approach, this introduces indirect data accesses, but delivers, if done correctly, excellent performance not only for flow in simple geometries but also in porous media like fixed-bed reactors or foams.

In this work we present a modified version of the two-grid one-step algorithm with non-temporal stores (OS-NT) and the AA-pattern algorithm. Both are augmented by a technique called RIA (reduced indirect addressing),

which can avoid the indirect access under certain conditions. This optimization is based on the idea of run length encoding and enables a reduction of the loop balance B_l. It also allows for partial vectorization, which is usually incompatible with indirect addressing unless the hardware supports efficient gather operations. RIA and partial vectorization were already implemented in the ILBDC code we employed for our earlier analysis [15]; here we describe them in due detail.

Beyond the chip level we focus on optimizations regarding large scale parallel execution. The partitioning becomes more and more important with rising communication overhead. The goal here is to reduce the communication volume and possibly the number of neighbors of all partitions without impacting the single-core performance. In case of ILBDC partitioning is performed by cutting the adjacency list in chunks of equal size. The quality of the resulting partitions depends on how the adjacency list was set up, i.e., which enumeration function was chosen (see [17] for details). Locality-preserving enumeration schemes such as space filling curves (SFC) or lexicographic sorting (LS) with small blocking factors result in compact partitions but have the disadvantage of poor single core performance, because they lead to short loops and non-vectorizable code. Better single core performance can be achieved with large blocking factors for LS, which in turn yields partitions of lower quality. The ideal solution here depends on the simulation geometry and requires experimenting with the parameters to find a good point where both requirements are met. We propose a two-stage method to tackle this problem: In the first step space filling curves are used for the adjacency list setup, which will generate high quality partitions. After the partitions have been assigned to the solver processes, they are re-enumerated by LS (without a blocking factor). This leads to a partitioning with low communication overhead and high single core performance. No manual experimentation, auto-tuning, etc. is required, and it is independent of the number of partitions.

The paper is structured as follows. In Sect. 2 we give a brief introduction to the lattice Boltzmann methods. The compute cluster used for performance measurements is described in Sect. 3. The principles of OS-NT and AA-pattern and their optimizations are discussed in Sect. 4. These optimizations are evaluated with ILBDC on a current Intel Haswell processor in Sect. 5. In Sect. 6 large-scale performance is evaluated. Finally we summarize the results in Sect. 7.

2 Lattice Boltzmann Methods

Lattice Boltzmann Methods are derived from the lattice Boltzmann equation (LBE) and result from a discretization of the velocity space and a numerical discretization of the spatial and time derivatives of the Boltzmann equation [10, 18]. The discretization model is denoted DdQq, where d represents the spatial dimension and q is the velocity discretization. Typical implementations are D2Q9, D3Q15, or D3Q19. Each lattice node comprises q particle distribution functions (PDF) f_i with $i = 0, \ldots, q-1$, which are the central element for computation. The LBE is written as

$$f_i(x + c_i \Delta t, t + \Delta t) = f_i(x,t) + \Omega(f_i(x,t), f_i(x,t)), \qquad i = 0, \ldots, q-1, \tag{1}$$

with the PDFs f_i, the position vector x, the velocity vector c_i, and the collision operator Ω. The right-hand side of (1) describes the collision of PDFs of the current time step, whereas on the left-hand side the propagation takes place, transferring the newly computed post-collision PDFs to the neighbor nodes. This is visualized in Fig. 1 for a D2Q9 discretization. In Fig. 1a the central node is about to be updated. The read and collided PDFs are then propagated as depicted in Fig. 1b. In this work the two-relaxation time (TRT) collision operator from GINZBURG et al. [2] is used.

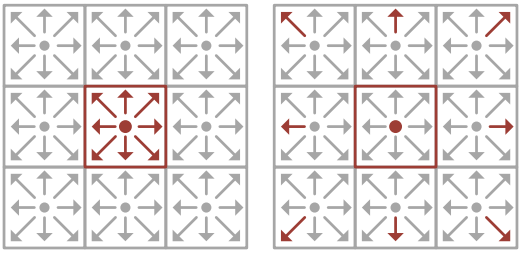

(a) Collision (b) Propagation

Figure (1) Collision (a) of the PDFs of the centering node and the following propagation (b) of the post-collision PDFs for a D2Q9 model.

Processor		
Type		Intel Xeon E5-2697 v3
Base freq.	[GHz]	2.6
Phys./SMT cores		14/28
ISA ext.		AVX, FMA3
Cache		
L1 (per core)	[KiB]	32, 8-ways
L2 (per core)	[KiB]	256, 8-ways
L3 (shared)[a]	[MiB]	2×17.5, 20-ways
TLB		
L1 4 KiB pages		64 4-ways
L1 2/4 MiB pages		32 4-ways
L2[b]		1024 8-ways

(a)

Compute node		
Sockets		2
Memory	[GiB]	4×16
Cluster-on-Die		enabled
NUMA LDs		4 (2 per CPU)
NUMA policy		default
Huge Pages	[MiB]	2

(b)

Table 1: Specification of the processor (a) and compute node (b) used for performance evaluation.

[a] In Cluster-on-Die mode two times seven cores share 17.5 MiB of L3 cache each.
[b] The L2 TLB is shared by 4 KiB and 2/4 MiB pages.

3 Test Bed

The SuperMUC Phase 2 cluster[1], located at the Leibniz Computing Center (LRZ) in Garching, Germany, was used as the benchmarking platform. This cluster is structured as six "islands" with 512 compute nodes each. The nodes are connected via InfiniBand FDR10, with a 4:1 over-subscribed fat tree among the islands. One node comprises two Intel Xeon E5-2697 v3 processors with 14 cores each. The CPUs are based on the Haswell microarchitecture and the cluster-on-die feature was enabled. In this mode the processor exhibits two ccNUMA locality domains (LD) instead of one, with seven cores per LD. Further details are listed in Tab. 1 or can be found in [6]. We used the Intel C/Fortran compiler 15.02 (with architecture-specific optimization parameters `-xCORE-AVX2` and `-fma`) and IBM MPI POE 1.4.

4 Optimized Propagation Step Implementations

The performance of highly optimized LBM kernels is limited by the memory bandwidth on current x86-based architectures. If the discretization model and the collision operator are not altered, i.e., if the number of FLOPs stays constant, performance can thus only be increased by reducing the loop balance B_l of a loop kernel.

In [16] we have shown that the two-grid one-step algorithm with non-temporal stores [12] (OS-NT), Bailey

[1] http://www.lrz.de/services/compute/supermuc/

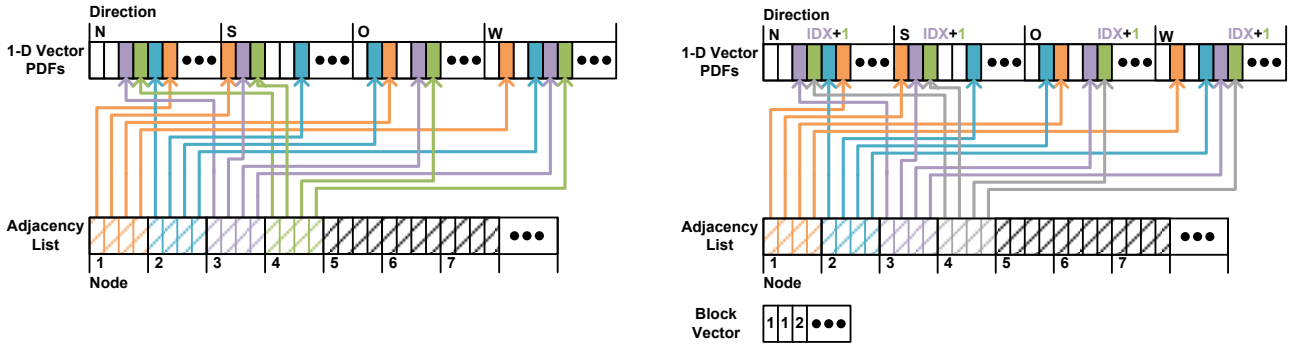

(a) Indirect Addressing

(b) Reduced Indirect Addressing

Figure 2: (a) Indirect addressing the PDF neighbors inside the 1-D vector via the adjacency list. (b) Indirect addressing of certain PDF neighbors can be avoided if the access pattern of the current node is the same as the previous one. This is the case for node 4.

et. al's AA-pattern [1], and Geier's Eso-Twist [7, 9] have the lowest loop balance. Here we concentrate on the first two, which are compatible with the code structure of ILBDC. In the following we assume a D3Q19 discretization model and use double-precision arithmetic.

Loop Balance

Independently of the propagation step during each update of a fluid node 19 PDFs must be loaded and stored, causing a data volume of

$$D_{\text{pdf}} = 2 \times 19 \times 8\ \text{B} = 304\ \text{B}. \tag{2}$$

When a PDF must be accessed indirectly, its index in the adjacency list must be obtained first (see Fig. 2a for a visualization of D2Q5). This causes an additional data volume of

$$D_{\text{idx}} = 18 \times 4\ \text{B} = 72\ \text{B}, \tag{3}$$

assuming a 4-byte index. Note that only 18 instead of 19 indirect accesses are needed, as the center PDF can always be directly addressed.

The OS-NT algorithm with loop splitting employs the "pull" scheme: First, PDFs from neighboring nodes are read via indirect addressing, collided, and the post-collision values are stored back to the local node via a direct access:

$$B_l(\text{OS-NT}) = \frac{D_{\text{pdf}} + D_{\text{idx}}}{1\ \text{FLUP}} = 376\ \frac{\text{B}}{\text{FLUP}}. \tag{4}$$

Loop splitting, which is required to reduce the number of concurrent non-temporal store streams, has no impact on the loop balance, as only data traffic between cores and memory is considered [12, 19] The AA pattern requires direct access in the even time step (D_{pdf}) and indirect access in the odd time step ($D_{\text{pdf}} + D_{\text{idx}}$). The average over both time steps is then

$$B_l(\text{AA}) = \frac{D_{\text{pdf}} + D_{\text{idx}}/2}{1\ \text{FLUP}} = 340\ \frac{\text{B}}{\text{FLUP}}. \tag{5}$$

4.1 Reduced Indirect Addressing

Reduced indirect addressing (RIA) is based on the idea that an indirect access is not necessary for a fluid node if the previous fluid node exhibits the same access pattern, but shifted by one. Then, the pointers to the

PDFs determined at the previous fluid node need only be incremented to get the current node's PDFs. As a consequence, the loop balance with RIA depends on how often the indirect access can be avoided and is thereby dependent on the structure of the simulation geometry and the chosen enumeration function to set up the adjacency list. The information about how many consecutive fluid nodes share the same access pattern (shifted by one) is stored as a run length coded (block) vector, as shown in Fig. 2b for a D2Q5 discretization. Each block vector access requires $D_{\text{block}} = 4$ B. One can calculate the minimum and maximum loop balance for both propagation patterns: The lower and upper loop balance bounds for the OS-NT algorithm with RIA (OS-NT-R) are

$$B_l(\text{OS-NT-R}) = \left[\frac{D_{\text{pdf}}}{1 \text{ FLUP}}, \frac{D_{\text{pdf}} + D_{\text{idx}} + D_{\text{block}}}{1 \text{ FLUP}} \right] = [304, 380] \frac{\text{B}}{\text{FLUP}}, \tag{6}$$

while for the AA pattern with RIA (AA-R) we get

$$B_l(\text{AA-R}) = \left[\frac{D_{\text{pdf}}}{1 \text{ FLUP}}, \frac{D_{\text{pdf}} + (D_{\text{idx}} + D_{\text{block}})/2}{1 \text{ FLUP}} \right] = [304, 342] \frac{\text{B}}{\text{FLUP}}. \tag{7}$$

4.2 Partial Vectorization

During the odd time step of AA the reading, writing, and collision of the PDFs is still performed in scalar mode. Via run length coding it is possible to *partially* SIMD-vectorize the fluid node updates if at least V nodes share the same access pattern, where V is the vector width (in case of AVX, $V = 4$). As the optimized LBM kernels are memory bound and this optimization does not alter the loop balance, no direct performance impact is expected when the kernel is in the memory bound regime, i.e., when sufficient cores are used to saturate the memory bandwidth. Hence, the loop balance for AA-RP, i.e., AA with RIA and partial vectorization, is still $B_l(\text{AA-R})$. However, we will show later that this optimization saturates the bandwidth with fewer cores, potentially saving energy.

4.3 Architectural Optimizations

Due to the high number of concurrent memory streams (38 with OS-NT and 19 with AA), D3Q19 bears the danger of cache thrashing. In order to avoid this problem we introduce a padding into the data vector, ensuring that different directions are always mapped on different sets of the L1 data cache. For OS-NT with 38 concurrent memory streams the L1 TLB with 32 entries for 2 MiB pages is not large enough, leading to cache thrashing. In contrast to previous Intel micro-architectures, a Haswell core provides an L2 TLB also for 2 MiB pages, which mitigates the effect. Although AA requires only 19 concurrent entries in the L1 TLB for 2 MiB pages, thrashing can still occur if more than four page table entries are mapped to the same set. Whether this happens depends on the number of ghost and fluid nodes. To generally avoid this problem we use an additional padding against TLB thrashing. Details about the impact of thrashing on performance can be found in [14].

5 Performance Results

In this section we analyze the performance characteristics within one ccNUMA LD, i.e., up to seven cores.

5.1 Benchmark Geometries

Two different simulation geometries are employed for benchmarking. The first is an empty channel with a quadratic cross section and dimensions of $500 \times 100 \times 100$ nodes containing around 4.8×10^6 fluid nodes. Due

Geometry	B_I(OS-NT)	B_I(OS-NT-R)	B_I(AA)	B_I(AA-R)
Channel	376	306	340	305
Fixed-bed reactor	376	333	340	319

Table 2: Loop balance (in B/FLUP) of the benchmark geometries for the OS-NT and AA-Pattern propagation step implementations with and without RIA.

Frequency	CNT-1A	CNT-19A	U-1A	U-19A
1.2 GHz	27.2	24.0	26.7	24.8
2.6 GHz	27.2	24.0	26.7	25.1

Table 3: Measured memory bandwidths (in GB/s) with different micro-benchmarks on seven cores of the benchmarked Haswell system, i. e. in one ccNUMA locality domain.

to the partial vectorization of the odd time step in AA, 98 % of the fluid nodes can be updated with SIMD instructions. The second geometry originates from a real-world application and represents a fixed-bed reactor of the same dimensions, but with only 2.1×10^6 fluid nodes. The reason for the low partial vectorizablility of only 57 % is the high porosity of this setup. The loop balance for both propagation step implementations is listed in Table 2.

5.2 Roofline Performance Model

We use the Roofline performance model [13] to calculate upper performance limits on the chip level, considering only the memory-bound case. Typically the result of the STREAM copy benchmark [8] is used as a bandwidth limit. Here the variant with non-temporal stores achieved the highest bandwidth, shown for one LD (seven cores) in Table 3. As shown in [3], the prediction of the Roofline model for LBM kernels can be improved if the streaming benchmark matches the access pattern of the application. Therefore we utilize two different micro-benchmarks for OS-NT and AA-pattern: CNT-19A has the same characteristics as OS-NT. 19 arrays are copied chunk-wise into a small temporary array, from which data is moved via nine loops, each copying two chunks utilizing non-temporal stores, to the final array. With U-19A the memory access pattern of the even time step of the AA-pattern is mimicked by updating 19 arrays concurrently. The measured bandwidths of both benchmarks are given in Table 3. For reference we also list the bandwidths attainable by copying (with non-temporal stores) and updating one array, named CNT-1A and U-1A, respectively. As already mentioned the saturated memory bandwidth on the Haswell system is rather independent of the core frequency. Consequently we see only a very small difference between the bandwidths at 1.2 GHz and 2.6 GHz in Table 3. The final predictions of the Roofline model can be found in Table 4. They are visible as horizontal bars in Fig. 4.

		OS-NT		**AA**	
		w/o RIA	w/ RIA	w/o RIA	w/ RIA
Channel					
Loop Balance	[B/FLUP]	376	306	340	305
P @ 1.2 GHz	[MFLUP/s]	63.8	78.4	72.9	81.3
P @ 2.6 GHz	[MFLUP/s]	63.8	78.4	73.8	82.2
Fixed-Bed Reactor					
Loop Balance	[B/FLUP]	376	333	340	319
P @ 1.2 GHz	[MFLUP/s]	63.8	72.0	72.9	77.8
P @ 2.6 GHz	[MFLUP/s]	63.8	72.0	73.8	78.8

Table 4: Performance prediction of the Roofline model when propagation step specific bandwidth micro-benchmarks are used.

port	instructions	ET	OTB	OTW
0	FMA, MUL	146	144	480
1	FMA, ADD	172	174	1080
2 & 3	LD	110	129	460
4	ST	72	114	360

Table 5: Distribution of the in-core execution time in cycles on the core execution ports for updating eight fluid nodes with AVX-vectorized AA-RP. Predictions are given for the even time step (ET) and the odd time step in the best case (OTB) and worst case (OTW).

5.3 ECM Performance Model

The Execution-Cache-Memory (ECM) performance model [4] considers each layer of the memory hierarchy separately and distinguishes between data transfers and arithmetic. For brevity we only model the AVX-vectorized AA-RB here. Within the ECM model cache line transfers play a major role. Thus the modeling is done at a granularity of eight fluid node updates, as (at least in the even time step) full cache lines are required. The in-core execution time predictions were determined via IACA [5] in throughput analysis mode, which assumes that full overlap of instructions is possible by out-of-order execution. The even time step (ET) could be analyzed without modifications. The kernel of the odd time step (OT) contains two conditional branches, which are responsible for deciding about indirect access and SIMD vectorization. IACA only considers branches from loop iterations, but not branches resulting from if conditions in the loop body. Therefore two kernels representing the best and worst cases of OT described in Sect. 4.1 were built. The best case (OTB) contains only direct accesses with vectorized execution, whereas in the worst case (OTW) only indirect access and scalar execution occurs. The results of the analysis are shown in Table 5.

For modeling the data paths the following bandwidths (in cycles per cache line, cy/cl) between adjacent cache/memory levels are assumed: 1 cy/cl between L1/L2, 2 cy/cl between L2/L3, 3.1 cy/cl at 1.2 GHz and 6.6 cy/cl at 2.6 GHz between L3/memory. The latter two values are calculated from the saturated memory bandwidth measurements (see above), while all others are documented by Intel.

To update eight fluid nodes during the even time step, 2×19 cache lines must be transferred. This is also the case for OTB. OTW requires additional traffic per fluid node: 8×4 B for the indirect access and 4 B for the block vector. For the work unit of eight fluid nodes this amounts to 4.5 cache lines. The execution diagram for the best and worst case is shown in Fig. 3.

ECM assumes linear performance scaling across cores until a bottleneck is reached. As the Haswell CPU has a scalable L3 cache, the memory bandwidth is the only potential bottleneck. The resulting prediction are depicted as green lines in Fig. 4.

5.4 Performance Results

The flow solver is purely MPI parallel. Each MPI process was pinned to a separate physical core, so the memory allocated by each process was ensured to reside inside the core's LD. Performance results are shown in Fig. 4.

The performance difference between the implementations at 1.2 GHz and 2.6 GHz is only around 10 %, independent of the simulation geometry. As with a lower frequency more cores are needed to saturate the memory bandwidth, this also holds true for the benchmarked propagation steps. Especially with the channel geometry and OS-NT(-R) as well as AA(-R), performance saturation already occurs with fewer than seven cores at the higher frequency setting. For AA-RP, where at both frequencies a saturation occurs at four cores already. This is a consequence of the partial vectorization, which already boosts the single core performance to over twice the level reached with AA and RIA alone.

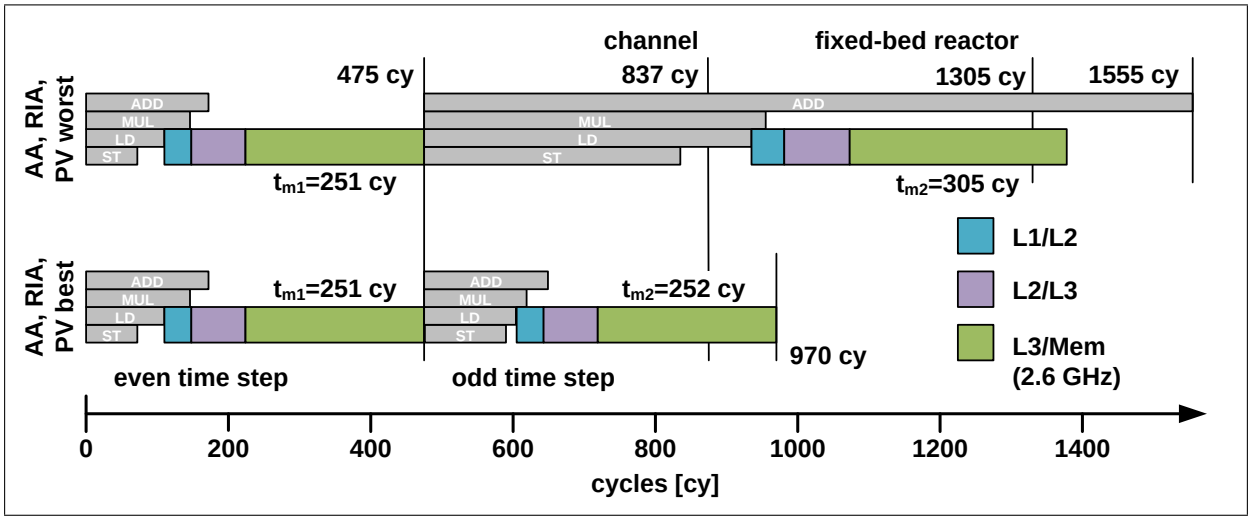

Figure 3: Update of eight fluid nodes on one Haswell core at 2.6 GHz modeled with the ECM model for the even and odd time step of AA with RIA for best and worst case.

The predictions of the Roofline model at 2.6 GHz are too pessimistic as OS-NT and also AA exceed the predicted performance. We think that the synthetic micro-benchmarks deliver a lower bandwidth than the application code, as they have no precautions against cache/TLB thrashing. According to the loop balance numbers in Table 4, OS-NT-R (AA-R) should reach a 20% (10%) higher performance for the channel geometry compared to the propagation step implementations without RIA, but only 15% (5%) improvement are actually achieved at 2.6 GHz. For the fixed-bed reactor the loop balances are only 10% (6%) lower with RIA than without for OS-NT (AA). Here both OS-NT-R and AA-RP see no benefit at 2.6 GHz.

The performance results for the propagation steps for the fixed-bed reactor, when all seven cores are utilized, are at most only 10% lower than with the channel geometry. In this case AA-RP does not increase performance that much at low core counts. The reason is that for this geometry only 57% of the fluid cells (on average) can be updated in a vectorized way during the odd time step. Hence the impact on the single core performance is lower.

The ECM model for the best case assumption matches well the AA-RB result on the channel geometry, because of its high vectorizablility. Here the model curves (green continuous lines) in Fig. 4 (a) and (b) nearly agree perfectly with the measured performance. Because of the low vectorizability the fixed-bed reactor performance should be close to the worst case assumption of ECM (green dashed lines). This is only the case for up to two/three cores, shown in Fig. 4 (c) and (d). With more cores the performance is lower than the prediction. This behavior is clearly caused by effects that are not part of the model; one candidate is misprediction of the conditional branches in the odd time step, caused by the small portions of fluid between the obstacles.

5.5 Energy Efficiency

For each benchmark run the power consumption of the core, uncore, and DRAM was measured via likwid [11]. Shown in Fig. 5 is the normalized energy to solution (NETS), i. e. measured power divided by measured performance in the unit J/MFLUP. The parameter on the curve is the number of cores utilized ranging from one to seven. OS-NT exhibits an interesting behaviour at 1.2 GHz: From three to five cores the NETS stays nearly constant while performance increases. With more cores NETS starts to decrease again. This effect is unclear and deserves a more detailed investigation. In all cases the NETS of AA is lower than of OS-NT, because of lower runtime resulting from the higher performance. The minimal NETS is reached, when performance saturates. If this point is reached when not all cores are utilized then each additional core will only increase the NETS

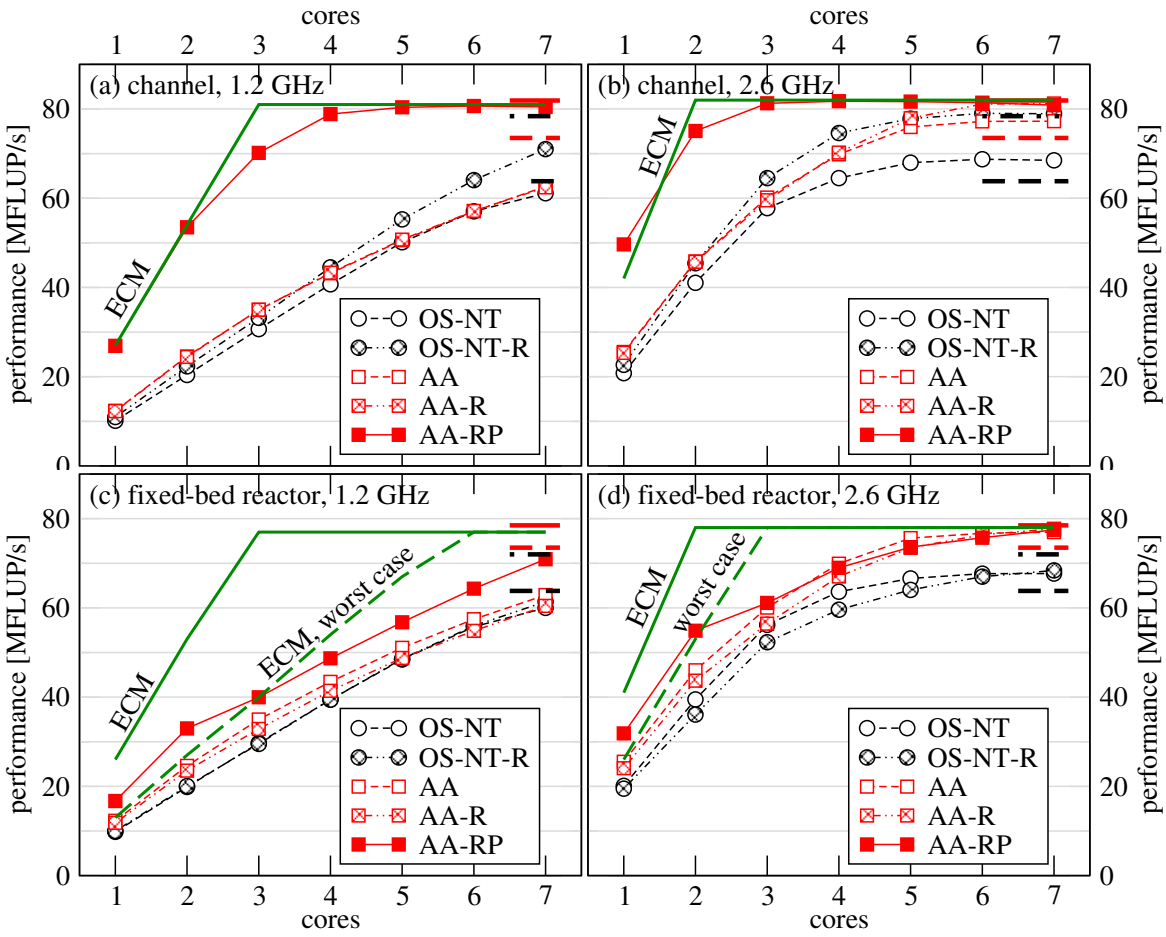

Figure 4: Performance of OS-NT and AA on one Haswell NUMA LD without and with RIA, respectively. Roofline limits (horizontal bars) for seven cores as well as the prediction of the ECM model for the best (green continuous line) and worst (green dashed line) case are shown as well.

but not performance, visible in Fig 5 (b) for AA-RP. The performance and NETS of AA-RP is always better than pure AA except for the fixed-bed reactor at 2.6 GHz. Here at saturated performance nearly no difference is visible. At most utilizing AA-RP in favor of AA reduces the NETS by around 15%.

5.6 Impact of Enumeration Function

For the measurements conducted in the previous subsections LS with blocking factor $B = 1$, i. e. actual without blocking, was used. Which enumeration function and parameters used for setting up the adjacency list has direct influence on the resulting solver performance. In Fig. 6 the loop balance B_l in cache and the performance for AAOptAvx is plotted for LS with different blocking and the space filling Hilbert curve. Hereby the channel geometry and one core of the Haswell system (2.6 GHz) was used. The loop balance in cache was determined via measuring the data traffic between L1 and L2 cache via likwid [11] for the even and odd time step separately. The in cache loop balance (black line) with $B_l = 304$ B/FLUP of the even time step is independent of the enumeration function, as here only accesses to the local node are required, which can always be performed via direct addressing. Only the odd time step is affected (red line). The plot shows that LS with and blocking factors ranging from $2 \leq B_l \leq 10$ the performance (blue dashed line) drops to half the performance achieved with $B = 1$. With this ordering cache lines during the odd time step are not efficiently used. They are evicted before all their containing PDFs have been used and therefore must be reloaded. The loop balance in this range

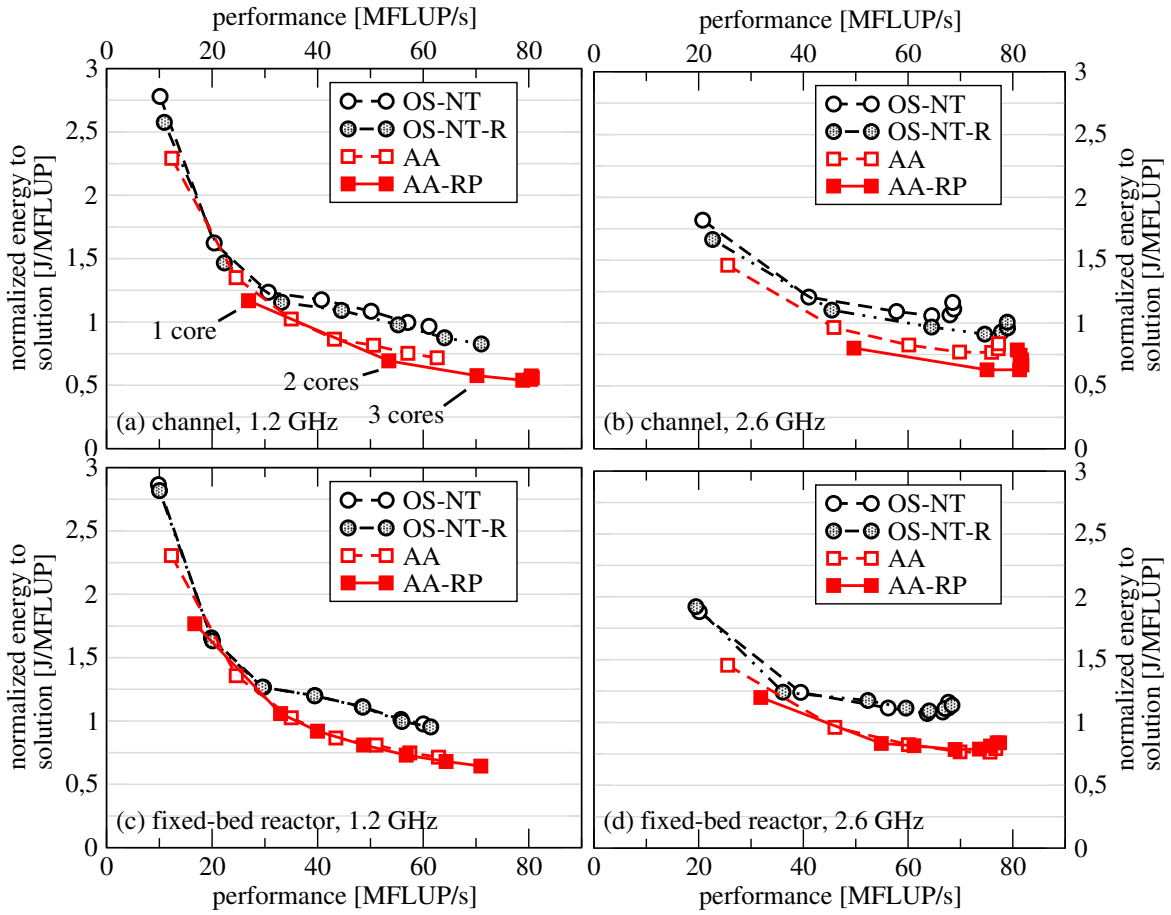

Figure 5: Performance vs. Normalized energy to solution for OS-NT and AA in the unoptimized and optimized variants, respectively. The parameter on the curve is the number of cores used, ranging from one to seven.

nearly doubles for the odd time step. This becomes the new bottleneck if too much time is spend with reloading cache lines. Furthermore the fraction of fluid nodes, which can be updated vectorized in the odd time step, drops between 0 % and 60 %. For larger blocking factors performance increases until with $B = 100$ the $B = 1$ performance is achieved again. The blocking factors $B = 1$ and $B = 100$ can in this case be considered equally as the diameter of the channel is 100 nodes and they generate the same ordering. Space filling curves like the Hilbert curve exhibit the same characteristics as LS with small blocking factors: high in cache loop balance and low vectorizablility of the odd time step as well as low performance. They suffer the same problem of inefficient cache line usage.

To reach a high single core performance LS with no blocking, i. e. $B = 1$, or with a high blocking factor is required. Small blocking factors or space filling curves tend to inefficient cache line usage and limit the vectorized update of fluid nodes which results in thereby a lower performance.

6 Large Scale Performance Aspects

For studying the strong scaling behaviour of ILBDC a fixed bed reactor with dimensions of $12500 \times 2500 \times 2500$ nodes is used. It contains around 34×10^9 fluid nodes, requiring 4.8 TiB for PDFs and 2.3 TiB for the adjacency list.

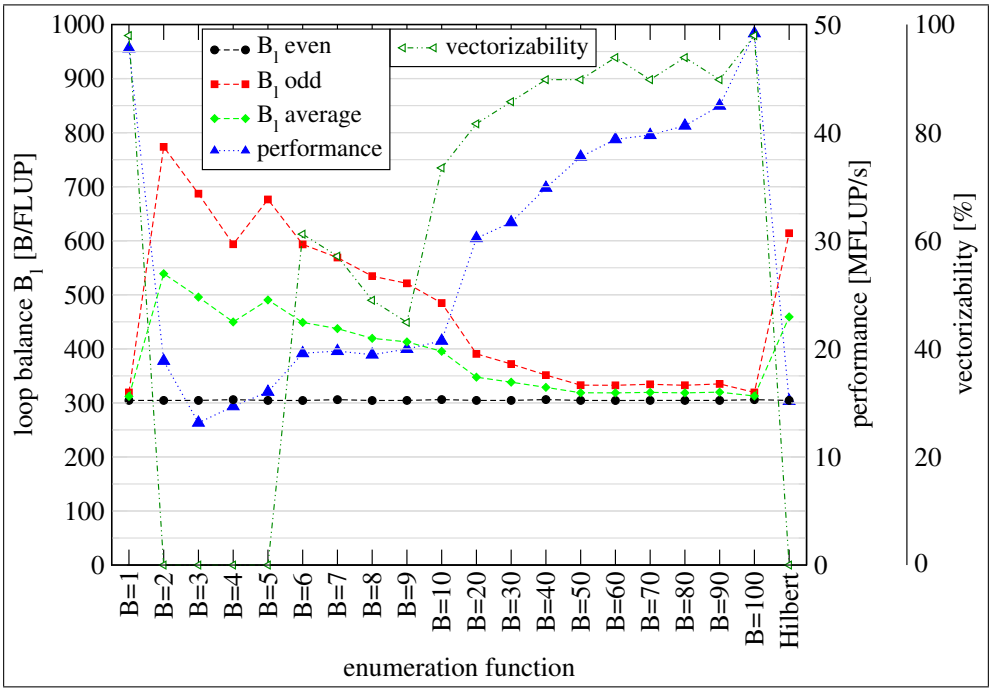

Figure 6: Performance and in-cache loop balance B_l of the even and odd time step of AA-RP on one core of the Haswell system. As enumeration functions LS with different blocking factors B as well as the space filling Hilbert curve are utilized. The loop balance was determined by measuring the data traffic between the L1 and L2 cache with likwid [11].

The performance results, when all physical cores of a compute node, i. e. PPN= 28, are used are shown in Fig. 7. When AA-RP and LS with blocking factor $B = 1$ at 2.6 GHz (right panel) is used a very low performance is reached. This is due to the high number of ghost PDFs, which must be exchanged between neighboring partitions. As with this blocking factor, the simulation domain is in principal cut perpendicular to the x direction. Each partition consists of a low number of slices which exhibit a high communication volume. At 256 nodes around 300 MiB per partition must be communicated, which drops to around 60 MiB for 1536 nodes. Through this high number of ghost cells per partition the memory of 256 nodes is exceeded, which is why there could be no measurement conducted. By choosing a higher blocking factor, e. g. $B = 100$, this can be mitigated. Here the communication volume per partition decreases to around and the overall performance increases. For 256 nodes only 8 MiB per partition must be communicated. This values decreases linearly to around 2 MiB at 1536 nodes. Using a space filling curve enumeration functions like the Hilbert curve leads also to partitions with a small surface. The communication volume per partitions is nearly the same as with LS and $B = 100$. But as already shown in Sect. 5.6 they only achieve a poor single core performance, which in turn leads to a low total performance. For reference also the performance of OS-NT-R with LS and $B = 100$ is plotted. Hence two grids are required only at 512 nodes enough memory was available. As the communication properties are the same as with AA-RP and LS with $B = 100$ the lower performance only results of the slower single core performance.

Reducing the clock speed to 1.2 GHz (left panel of Fig. 7) only decreases the performance marginally when AA-RP and LS with $B = 100$ is used. Instead for AA-RP with Hilbert and OS-NT-R with LS and $B = 100$ a decrease of 20 % and 30 % is shown, respectively. This stems to a large part from the lower single node performance.

At first sight super linear speedup is visible from the first to the second point on each curve. This is actually not the case, as the first points suffer from NUMA issues. Allocation of the lattices and adjacency lists is not synchronized. If a process' local NUMA LD has not enough free memory, when the lattice and the list is placed, because other processes have not already deallocated their temporary structures needed for setup, parts of other

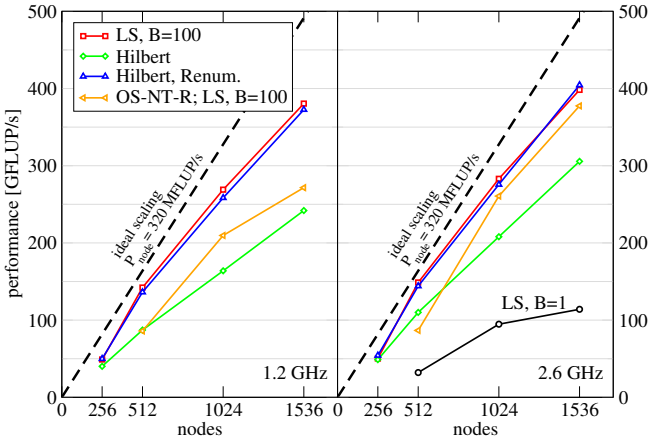

Figure 7: Performance of AA-RP with PPN= 28 on the SuperMUC Phase 2 cluster.

LDs are used. This decreases the single node performance as the bandwidth to a remote NUMA LD is lower than to the local one.

A high performance for large scale simulations depends on the choice of the right enumeration function and their parameter(s). On the one hand a high blocking factor is required to achieve a good single core performance, whereas on the other hand a small blocking factor (or an SFC) generate partitions with small surfaces. The goal is here to find a mechanism, which fulfills both conditions.

6.1 Renumbering

To achieve this goal a two step process is used. In the first step an SFC, like the Hilbert curve, is initially used as enumeration function. This ensures, that during partitioning the communication volume of each partition is minimized. In the second step the fluid nodes assigned to a solver process are renumbered. Here LS with blocking factor $B = 1$ can be chosen, for a high single core performance. The performance achieved is shown in Fig. 7 as "Hilbert renum." With this method nearly the same performance as with LS and the manually determined good parameter $B = 100$ is achieved.

6.2 Reducing the Number of Processes per Node

With AA-RP only three to four cores are required to saturate the performance, as shown in Sect. 5.4. This is also the case in large scale. In Fig. 8 the performance of AA-RP with Hilbert renumbered is shown then the number of processes per node (PPN) is reduced from 28 over 24 and 20 down to 16. As expected no performance impact is observed. Also with 1.2 GHz this behaviour is the same. Here the measurements from Fig. 4 (c) would indicate otherwise, as actually all cores are necessary to saturate performance. This behaviour is unclear and requires a more detailed elaboration.

6.3 Energy efficiency

At 1.2 GHz AA-RP nearly reaches the performance of 2.6 GHz also in the large scale case, if the correct enumeration function is chosen. Here energy savings of 40 % can be achieved with the Hilbert renumbered enumeration function, as shown in Fig. 9. Reducing the number of processes on each node, i. e. reducing the number of active cores, from PPN= 28 down to PPN= 16 only a small fraction of energy is saved. This is the

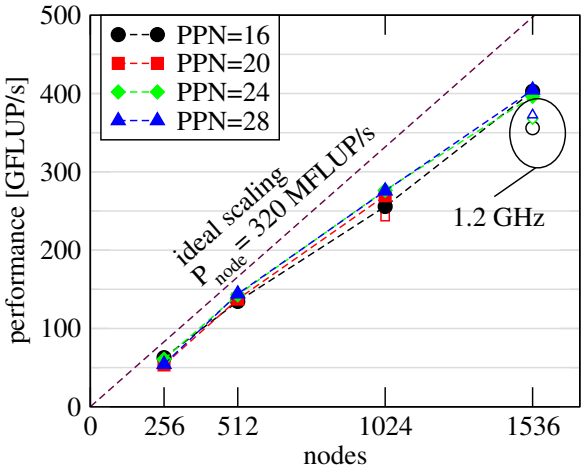

Figure 8: Performance for AA-RB with Hilbert renumbered on the SuperMUC Phase 2 cluster at 2.6 GHz for different number of processes per nodes (PPN).

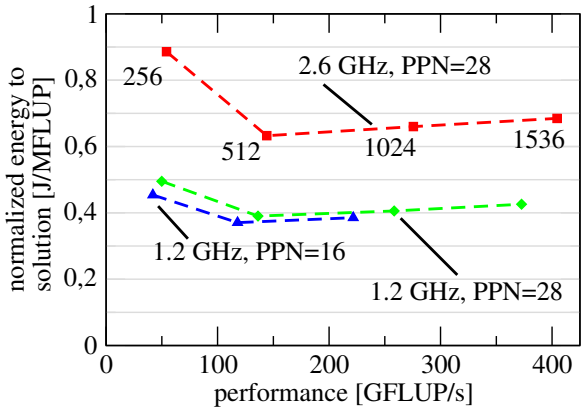

Figure 9: Performance in relation to the normalized energy to solution (NETS) for AA-RB with Hilbert renumbered on the SuperMUC Phase 2 cluster.

same behaviour as in one NUMA LD for the fixed-bed reactor at 1.2 GHz as shown in Fig. 5 (c). Note that the value for 1536 nodes in Fig. 9 for 1.2 GHz and PPN= 16 is missing.

7 Conclusion

Reduced indirect addressing (RIA) avoids for certain nodes the indirect access and thereby reduces the data traffic per node update, i. e. the loop balance. The implementation of RIA for the propagation step OS-NT showed that depending on the simulation geometry a performance increase of 15% is observable. For the AA-pattern with RIA no significant performance increase is visible, as savings in the loop balance are less than for OS-NT with RIA. AA-pattern with RIA enables the partial vectorization of the odd time step, which has more impact on performance. Hereby the loop balance is not altered and thus the saturated performance is not higher than with only RIA enabled, but saturation can be reached with less cores. With lower frequencies like 1.2 GHz instead of 2.6 GHz no saturation occurs. At this point partial vectorization increases the overall performance by up to 30% depending on the simulation geometry. The savings in the energy consumption (up to 30 %) are higher with simple structured simulation geometries as there the benefit from partial vectorization is higher.

In the large scale case choosing the right enumeration function is crucial. With a simple two step method we

avoid the manual search for good parameters, which depend on the simulation geometry, and still reach the same performance. Also here are the single core optimizations visible, which allow to run at lower frequencies and use less cores. Here up to 40 % of energy savings are possible.

We would like to thank the Leibniz Computing Center (LRZ) in Garching, Germany for the great support during conducting the benchmarks. This work was financially supported by KONWIHR III and partially supported by BMBF under grant No. 01IH08003A (project SKALB).

References

[1] P. Bailey, J. Myre, S. Walsh, D. Lilja, and M. Saar. Accelerating lattice Boltzmann fluid flow simulations using graphics processors. In *International Conference on Parallel Processing 2009 (ICPP'09)*, pages 550–557, Sept 2009. doi:10.1109/ICPP.2009.38.

[2] I. Ginzburg, F. Verhaeghe, and D. d'Humières. Two-relaxation-time lattice Boltzmann scheme: About parametrization, velocity, pressure and mixed boundary conditions. *Commun. Comput. Phys.*, 3(2):427–428, 2008.

[3] J. Habich. *A Performance Engineering Process for Developing High Performance Lattice Boltzmann Implementations*. PhD thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg, 2015.

[4] G. Hager, J. Treibig, J. Habich, and G. Wellein. Exploring performance and power properties of modern multicore chips via simple machine models. *Concurrency and Computation: Practice and Experience*, 2014. doi:10.1002/cpe.3180.

[5] Intel Corp. Intel Architecture Code Analyzer. http://software.intel.com/en-us/articles/intel-architecture-code-analyzer, 2012. Version: 2.0.1.

[6] Intel Corp. Intel64 and IA-32 Architectures Optimization Reference Manual. http://www.intel.com/content/dam/doc/manual/64-ia-32-architectures-optimization-manual.pdf, 2015. Version: September 2015.

[7] J. Linxweiler. *Ein integrierter Softwareansatz zur interaktiven Exploration und Steuerung von Strömungssimulationen auf Many-Core-Architekturen*. PhD thesis, Fakultät Architektur, Bauingenieurwesen und Umweltwissenschaften, TU-Braunschweig, 2011. (in German).

[8] J. D. McCalpin. Memory bandwidth and machine balance in current high performance computers. *IEEE Computer Society Technical Committee on Computer Architecture (TCCA) Newsletter*, pages 19–25, Dec. 1995.

[9] M. Schönherr, M. Geier, and M. Krafczyk. 3D GPGPU LBM implementation on non-uniform grids. In *International Conference on Parallel Computational Fluid Dynamics 2011 (Parallel CFD 2011)*, May 2011. http://parcfd2011.bsc.es/sites/default/files/abstracts/id121-schonherr.pdf.

[10] S. Succi. *The Lattice Boltzmann Equation: For Fluid Dynamics and Beyond*. Numerical Mathematics and Scientific Computation. Oxford University Press on Demand, 2001.

[11] J. Treibig. LIKWID performance tools. http://code.google.com/p/likwid.

[12] G. Wellein, T. Zeiser, S. Donath, and G. Hager. On the single processor performance of simple lattice Boltzmann kernels. *Computers & Fluids*, 35:910–919, 2006. doi:10.1016/j.compfluid.2005.02.008.

[13] S. Williams, A. Waterman, and D. Patterson. Roofline: an insightful visual performance model for multi-

core architectures. *Commun. ACM*, 52(4):65–76, Apr 2009. doi:10.1145/1498765.1498785.

[14] M. Wittmann. *Hochparallele Implementierung von effizienten Lattice-Boltzmann-Verfahren für komplexe Geometrien*. PhD thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg, 2016. to appear, in German.

[15] M. Wittmann, G. Hager, T. Zeiser, J. Treibig, and G. Wellein. Chip-level and multi-node analysis of energy-optimized lattice Boltzmann CFD simulations. *Concurrency and Computation: Practice and Experience*, 2015. doi:10.1002/cpe.3489.

[16] M. Wittmann, T. Zeiser, G. Hager, and G. Wellein. Comparison of different propagation steps for lattice Boltzmann methods. *Computers & Mathematics with Applications*, 65(6):924–935, 2013. doi:10.1016/j.camwa.2012.05.002.

[17] M. Wittmann, T. Zeiser, G. Hager, and G. Wellein. Domain decomposition and locality optimization for large-scale lattice Boltzmann simulations. *Computer & Fluids*, 80:283–289, 2013. doi:10.1016/j.compfluid.2012.02.007.

[18] D. Wolf-Gladrow. *Lattice-Gas Cellular Automata and Lattice Boltzmann Models: An Introduction*. Number 1725. Springer, 2000.

[19] T. Zeiser, G. Hager, and G. Wellein. Benchmark analysis and application results for lattice Boltzmann simulations on NEC SX vector and Intel Nehalem systems. *Parallel Processing Letters*, 19(4):491–511, 2009. doi:10.1142/S0129626409000389.